

A Comparison and Evaluation of Map Construction Algorithms

Mahmuda Ahmed · Sophia Karagiorgou ·
Dieter Pfoser · Carola Wenk

Received: date / Accepted: date

Abstract Map construction methods automatically produce and/or update road network datasets using vehicle tracking data. Enabled by the ubiquitous generation of geo-referenced tracking data, there has been a recent surge in map construction algorithms coming from different computer science domains. A cross-comparison of the various algorithms is still very rare, since (i) algorithms and constructed maps are generally not publicly available and (ii) there is no standard approach to assess the result quality, given the lack of benchmark data and quantitative evaluation methods. This work represents a first comprehensive attempt to benchmark map construction algorithms. We provide an evaluation and comparison of seven algorithms using four datasets and four different evaluation measures. In addition to this comprehensive comparison, we make our datasets, source code of map construction algorithms and evaluation measures publicly available on mapconstruction.org. This site has been established as a repository for map construction data and algorithms and we invite other researchers to contribute by uploading code and benchmark data supporting their contributions to map construction algorithms.

Mahmuda Ahmed
University of Texas at San Antonio
San Antonio, TX, USA
E-mail: mahmed@cs.utsa.edu

Sophia Karagiorgou
National Technical
University of Athens, Greece
E-mail: sokaragi@mail.ntua.gr

Dieter Pfoser
George Mason University
Fairfax, VA, USA
E-mail: dpfoser@gmu.edu

Carola Wenk
Tulane University
New Orleans, LA, USA
E-mail: cwenk@tulane.edu

Keywords tracking data · map construction · quality measures · algorithms · performance

1 Introduction

Street maps and transportation networks are of fundamental importance in a wealth of applications. In the past, the production of street maps required expensive field surveying and labor-intensive postprocessing. Proprietary data vendors such as Navteq, TeleAtlas and Google therefore dominated the market. Over the last years, Volunteered Geographic Information (VGI) (23) efforts such as OpenStreetMap (OSM) (25; 32) have complemented commercial map datasets. They provide map coverage especially in areas, which are of less commercial interest. VGI efforts however still require dedicated users to author maps using specialized software tools. Lately, on the other hand, the commoditization of GPS technology and integration in mobile phones coupled with the advent of low-cost fleet management and positioning software has triggered the generation of vast amounts of tracking data. As a size indicator one can consider the contribution of tracking data in OpenStreetMap, which is steadily increasing in size and currently amounts to 2.6 trillion points (31). Besides the use of such data in traffic assessment and forecasting (18), i.e., map-matching vehicle trajectories to road networks to obtain travel times (9), there has been a recent surge of actual *map construction algorithms* that derive not only travel time attributes but actual road network geometries from tracking data, e.g., (1; 2; 10; 4; 7; 8; 11; 12; 13; 22; 16; 17; 24; 19; 26; 27; 29; 36; 40; 38; 37; 42). Among those, only a few algorithms give theoretical quality guarantees (1; 4; 13). An example of a constructed road network is given in Figure 1, which shows (a) the vehicle trajectories collected for Berlin and (b) the respective constructed road network using the algorithm of (27) with an OpenStreetMap background map.




Fig. 1: Vehicle tracking data vs actual road network.

A major challenge in the research community is to compare the performance and to evaluate the quality of the various map construction algorithms. Visual inspection remains the most common evaluation approach throughout the literature and only a few recent papers incorporate quantitative distance measures (3; 7; 8; 27; 29). However, the cross-comparison of different algorithms remains

rare, since algorithms and constructed maps are generally not publicly available. Also, there is a lack of benchmark data, and the quantitative evaluation with suitable distance measures is in its infancy. A cultural shift has recently been triggered by Biagioni and Eriksson (7), who perform a quantitative evaluation of several map construction algorithms. They have made their implementations of the algorithms as well as their dataset publicly available. The present work *significantly expands these benchmarking efforts to provide an evaluation and comparison of more map construction algorithms on more diverse datasets using various quality measures suitable for different applications*. Such an effort can only be sustained in a culture of sharing that makes data, methods and source code publicly available.

In this work, we evaluate and compare seven map construction algorithms using four benchmark tracking datasets and four different distance measures. The *algorithms* we compare are the state-of-the-art of the last years. The algorithms we evaluate include the recent algorithms by Ahmed and Wenk (4), by Karagiorgou and Pfoser (27), and by Ge et al. (22), in addition to the algorithms by Cao and Krumm (11), Davies et al. (16), Edelkamp and Schrödl (17), and Biagioni and Eriksson (8). Among those (11), (16) and (17) were previously compared by Biagioni and Eriksson (7).

The *four distance measures used to assess the constructed map quality* comprise two novel distance measures that have not been used for comparative evaluations of map constructions before and that work with unmodified and unbiased ground-truth maps: the Directed Hausdorff distance (6) and the path-based distance measure presented by Ahmed et al. (3). We also use a distance measure based on shortest paths by Karagiorgou and Pfoser (27) and the graph-sampling-based map comparison by Biagioni and Eriksson (8).

The *tracking datasets* include the *Chicago* dataset provided by Biagioni and Eriksson (7; 8), and three additional tracking datasets: two from *Athens*, Greece and one from *Berlin*, Germany (see detail in Section 4). They are available together with unmodified ground-truth maps obtained from OpenStreetMap. We use different datasets because they cover diverse roads (i.e. highways, secondary roads), different sampling rates and different scale.

In addition to providing the largest comprehensive comparison of map construction algorithms, we make our three new benchmark datasets, the map construction algorithms and outputs by Ahmed and Wenk (4) and by Karagiorgou and Pfoser (27), as well as the metric code for computing the three distance measures: the Directed Hausdorff distance (6), the path-based distance (3) and shortest-path based measure (27) publicly available on the internet at mapconstruction.org. We have established this Web site as a repository for map construction data and algorithms, and we invite other researchers to contribute by uploading code and benchmark data supporting their map construction algorithms. We expect that such a central repository will encourage a culture of sharing and will enable the development of improved map construction algorithms.

Our main goal with this work is to provide a common platform to do comparative analysis of map construction algorithms. As different distance measures capture different features of a constructed map, it is hard to combine them into a single score and rank the algorithms based on that. Also, which algorithm is the best highly depends on the quality of the input data and for what purpose the map will be used. For example, for the *Chicago* dataset the KDE-based algorithm by Davies et al. (16) generates a very good-quality map in terms of spatial dis-

tance to the ground-truth map (captured using path-based and Directed Hausdorff distance), but if the user is interested in maps with good coverage (captured by shortest-path based and graph-sampling based distance measure) this algorithm will not be the best choice as it ignores tracks in sparse areas as outliers/noise. So, we leave it to the user to pick the distance measure that suits his/her needs best.

The outline of the paper is as follows. Section 2 surveys map construction algorithms by introducing categories for types of algorithms and giving more details on the algorithms that we will use in our evaluation. Section 3 discusses quality measures that will allow us to assess the quality of the constructed maps. The tracking datasets that we provide for evaluation purposes are briefly discussed in Section 4. The datasets are available for download and also include the respective ground-truth map data. A comprehensive performance study comparing the various algorithms across datasets is given in Section 5. Finally, Section 6 provides conclusions and directions for future work.

2 Map Construction Algorithms

We assume that the input is given as a set of *tracks*, where each track is a sequence of *measurements*. Each measurement consists of a point (latitude/longitude or (x, y) -coordinates after suitable projection), a time stamp, and optionally additional information such as vehicle heading or speed. The desired output is to construct a *street map*. There are many possible models for street maps, mostly depending on the desired application and granularity. For example, an intersection can be modeled as a single vertex embedded as a point in the plane, or it could be a set of vertices, possibly annotated with turn restrictions, or it could be a region. And an edge can be modeled just as an abstract connection between vertices, as a curve embedded in the plane, as a set of curves to model multiple lanes, and an edge might be directed to model one-way streets. We will focus on the most basic model of a street map as an undirected geometric graph, where each vertex is embedded as a point in the plane and each edge is a polygonal curve that connects two vertices. All map construction algorithms in the literature follow this basic model; in addition, some algorithms enhance this basic model with some additional information (such as directions, turn restrictions, number of lanes), which often is computed in an additional post-processing step.

2.1 Related Work

There exist several different approaches in the literature for constructing street maps from tracking data. These can be organized into the following categories: Point clustering (this includes k -means algorithms and Kernel Density Estimation (KDE) as described in Biagioni and Eriksson (8)), incremental track insertion, and intersection linking.

2.1.1 Point Clustering

Algorithms in this general category assume the input consists of a set of points which are then clustered in various different ways to obtain street segments which

finally connect to a street map. The input point set either comprises the set of all raw input measurements, or a dense sample of all input tracks. Here, the input tracks are assumed to be continuous curves obtained from interpolating (usually piecewise-linearly) between measurements.

One approach employs the k -means algorithm to cluster the input point set, using distance measures (e.g., Euclidean distance) and possibly also vehicle heading of the measurement, as a condition to introduce seeds at fixed distances along a path. This includes Edelkamp et al. (17), who develop algorithms for road segmentation, map-matching, and lane clustering. In (36) this approach was used to refine an existing map rather than building it entirely from scratch. In (24), Guo et al. make use of statistical analysis of GPS tracks. However, the assumptions on the GPS data for symmetric 2D Gaussian distribution becomes unrealistic, especially in error-prone environments. Worrall et al. (40) emphasize compression of the input tracks to infer a digitized road map and present their results only for small datasets. They are mostly concerned about topological elements and not on connected way points. Similarly, Jang et al. (26) proposed a system of map construction with less than ten traces and presented in a very small scale and without any reference to the data features (i.e., sampling rate, GPS error). Agamennoni et al. (2) presented a machine-learning method to consistently build a representation of the map mostly in dynamic environments such as open-pit mines. Liu et al. (29) first cluster line segments based on proximity and direction, and then use the resulting point clusters and fit polylines to them, to extract road segments.

Another approach employs KDE methods to first transform the input point set to a density-based discretized image. Most of the KDE algorithms function well either when the data is frequently sampled (i.e., once per second) (12), or when there is a lot of data redundancy (8; 38; 37; 16). A similar approach to (8) is presented in Liu et al. (29). Generally, KDE algorithms have a hard time to overcome the problem of noisy samples when they accumulate in an area. Recently, Wang et al. (39) addressed the problem of map updates by applying their approach to OpenStreetMap data using a KDE-based approach.

In the computational geometry community, map construction algorithms have been proposed that cluster the input points using local neighborhood properties by employing Voronoi diagrams, Delaunay triangulations (13; 22), or other neighborhood complexes such as the Vietoris-Rips complex (1). All these algorithms assume a densely sampled input point set, and provide theoretical quality guarantees for the constructed output map, under certain assumptions on the underlying street map and the input tracks. Aanjaneya et al. (1) view street maps as metric graphs, and they focus on computing the combinatorial structure by computing an almost isometric space with lower complexity, but they do not compute an explicit embedding of vertices and edges. Chen et al. (13) focus on detecting “good” street portions in the road network and connecting them subsequently. The theoretical quality guarantees, however, assume a dense point sample coverage and error bounds, and make assumptions on the road geometry.

2.1.2 Incremental Track Insertion

Algorithms in this category construct a street map by incrementally inserting tracks into an initially empty map (30), often making use of map-matching ideas (33). Distance measures and vehicle headings are also used to perform additions

and deletions during the incremental construction of the map. One of the first algorithms in this category (35) clusters the tracks merely to refine an existing road network and not to compute it from scratch. Cao and Krumm (11) first introduce a clarification step in which they modify the input tracks by applying physical attraction to group similar input tracks together. Then they incrementally insert each track by using local criteria such as distance and direction. Bruntrup et al. (10) propose a spatial-clustering based algorithm that requires high quality tracking data (sampling rate and positional accuracy). The work in (42) discusses a road network update algorithm based on spatial similarity. It uses a method similar to GPS trace merging to continuously refine existing road maps. Ahmed and Wenk (4) present an incremental method that employs the Fréchet distance to partially match the tracks to the map.

2.1.3 Intersection Linking

While related to point clustering, the intersection linking approach is to first detect the intersection vertices of the street map, and in a second step link those intersections together by identifying suitable street segments. Fathi and Krumm (19) provide an approach that detects intersections by using a prototypical detector trained on ground truth data from an existing map. While a road network is finally derived, their approach works best for well aligned road networks and it uses frequently sampled data of 1s or 5s. The method by Karagiorgou and Pfoser (27) relies on detecting changes in the direction of movement to infer intersection nodes, and then “bundling” the trajectories around them to create the network edges.

2.2 Compared Algorithms

Here we give some more details on the map construction algorithms that we compare in Section 5. The algorithms categories are also provided in Table 1.

Algorithm	Point Clustering	Incremental Track Insertion	Intersection Linking
Ahmed and Wenk (4)		✓	
Biagioni and Eriksson (8)	✓		
Cao and Krumm (11)		✓	
Davies et al. (16)	✓		
Edelkamp and Schrödl (17)	✓		
Ge et al. (22)	✓		
Karagiorgou and Pfoser (27)			✓

Table 1: Algorithms categories.

2.2.1 Ahmed and Wenk (4)

The algorithm by Ahmed and Wenk (4) is a simple and practical incremental track insertion algorithm. The insertion of one track proceeds in three steps. The first

step performs a partial map-matching of the track to the partially constructed map in order to identify matched portions and unmatched portions. This partial map-matching is based on a variant of the Fréchet distance. In the second step, the unmatched portions of the track are then inserted into the partially constructed map, which requires creating new vertices and creating and splitting edges. In a third step, the already existing edges in the map that are covered by the matched portions of the trajectory, are updated using a minimum-link algorithm to compute a new representative edge. This last step is only needed to provide a guaranteed bound on the complexity of the output map; in the implementation of this algorithm that we use in Section 5, this last step has been omitted. Ahmed and Wenk also give theoretical quality guarantees for the output map computed by their algorithm, which include a one-to-one correspondence between well-separated “good” portions of the underlying map and the output map, with a guaranteed Fréchet distance between those portions.

2.2.2 Biagioni and Eriksson (8)

Biagioni and Eriksson (8) describe a point clustering-based algorithm that uses KDE methods. Their algorithm proceeds in using KDE to compute a skeleton map and they annotate the map by performing a map-matching pass of the input tracks with the skeleton map.

2.2.3 Cao and Krumm (11)

This incremental track insertion approach proceeds in two stages. In the first stage, simulation of physical attraction is used to modify the input tracks to group portions of the tracks that are similar together. This results in a cleaner data set in which track clusters are more pronounced and different lanes are more separated. Then, this much cleaner data is used as the input for a fairly simple incremental track insertion algorithm. This algorithm makes local decisions based on distance and direction to insert an edge one vertex and either merge the vertex into an existing edge, or add a new edge and vertex.

2.2.4 Davies et al. (16)

This is a classical KDE-based map construction algorithm. It first computes for each grid cell the density of tracks that pass through it. Then it computes the contour of the resulting bit map, and then it uses the Voronoi diagram of the contour to compute a center line representation, followed by additional cleanup.

2.2.5 Edelkamp and Schrödl (17)

Edelkamp and Schrödl (17) were the first who proposed a map construction approach based on the k -means method. Their point clustering algorithm creates road segments based on tracking data, represents the center line of the road using a fitted spline and performs lane finding. The lanes are found by clustering tracks based on their distance from the road center line.

2.2.6 Ge et al. (22)

This algorithm is a point clustering approach that applies topological tools to extract the underlying graph structure. They assume a densely sampled point set and their algorithm only needs a distance matrix or proximity graph of the point set as input. They use the Vietoris-Rips complex, or a more general proximity graph, to obtain a first layer of organization for the input point set. They also define a function f on the graph which captures the geodesic distance to a base point. Then they employ the Reeb graph to model the connected components of the level set $f^{-1}(a)$ for varying a . They provide runtime guarantees as well as partial quality guarantees for correspondences of cycles. An embedding for the edges is then obtained by using a principal curve algorithm (28) that fits a curve to the points contributing to the edge.

2.2.7 Karagiorgou and Pfoser (27)

This intersection-linking map construction algorithm is a heuristic approach that “bundles” trajectories around intersection nodes. Hence also the name *TraceBundle* algorithm. The main contribution of *TraceBundle* is as to how it derives intersection nodes. The basic heuristic relies on detecting *changes in movement* and then clustering “similar” nodes. Change, observed as a change in direction and speed, is considered a turn indicator. Clustering these turns based on (i) spatial proximity and (ii) turn type results in *turn clusters*. The centroid location of each of these turn clusters represents then an *intersection node*. Connecting the trajectories to intersection nodes and compacting them, they then generate links and consequently the entire geometry of the road network. The essential steps of the *TraceBundle* algorithm are as follows. 1. *Turn samples* - given a trajectory, each node (position sample) at which a significant change in direction and speed (parameters) occurs becomes a turn sample. 2. *Turn clusters* - clustering turn samples based on (i) proximity (*static parameter*) and (ii) a *turn model*. 3. *Intersection nodes* - centroid of turn clusters. 4. *Connecting intersection nodes* - using constituting turn samples, connect trajectories to respective intersection nodes. 5. *Compacting links* - merge connecting trajectory portions between intersection nodes to generate links.

3 Quality Measures for Map Comparison

There are two key ingredients for evaluating the quality of a constructed map: (1) the availability of an adequate ground-truth map G as part of the benchmark data, and (2) a quality measure used to evaluate the similarity between the constructed map C and the ground-truth map G .

There are essentially two cases of what can be considered as a ground-truth map G . Ideally, G is the underlying map consisting of all streets, and only those streets, that have been traversed by the entities that generated the set of input tracks. If such a G was available, then a suitable quality measure would compare C to all of G and the ideal would be for C to equal G . However, in practice, it is hard to obtain an unbiased ground-truth map that exactly corresponds to the coverage of the tracking data. This non-trivial task has been addressed in the past either by pruning the ground-truth either manually, by proximity to the tracking

data, or by map-matching the tracking data to the map (7; 8; 27; 29). Clearly all these approaches introduce an undesired bias.

Actually, it is much easier to obtain a ground-truth map that covers a superset of all the streets covered by the input tracks, e.g., street maps taken by proprietary vendors or OpenStreetMap. Therefore, if G is a superset, then the quality measure attempts to *partially* match C to G . Of course, another possible scenario is that C contains additional streets that are not present in either variation of G .

3.1 Related Work

In the graph theory literature, there are various distance measures for comparing two abstract graphs, that do not necessarily have a geometric embedding (15; 34; 21). Most closely related to street map comparison are the subgraph isomorphism problem and the maximum common isomorphic subgraph problem, both of which are NP-complete. These, however, rely on one-to-one mappings of graphs or subgraphs, and they do not take any geometric embedding into account. Graph edit distance (20; 41) is a way to allow noise by seeking a sequence of edit operations to transform one graph into the other, however it is NP-hard as well. Cheong et al. (14) consider a graph edit distance for geometric graphs (embedded in two different coordinate systems, however), and also show that it is NP-hard to compute.

For comparing street maps, distance measures based on *point sets* and distance measures based on *sets of paths* have been proposed. Point set-based distance measures treat each graph as the set of points in the plane that is covered by all its vertices and edges. The idea is then to compute a distance between the two point sets. A straightforward distance measure for point sets are the directed and undirected Hausdorff distances (6). The main drawback of such an approach is that it does not use the topological structure of the graph. Biagioni and Eriksson (7; 29), use two distance measures that essentially both use a variant of a partial one-to-one bottleneck matching that is based on sampling both graphs densely.

The two distance measures compare the total number of matched sample points to the total number of sample points in the graph, thus providing a measure of how much of the graph has been matched. They do require though to have as input a ground-truth graph that closely resembles the underlying map and not a superset.

For path-based distance measures on the other hand, the underlying idea is to represent the graphs by sets of paths, and then define a distance measure based on distances between the paths. This captures some of the topological information in the graphs, and paths are of importance for street maps in particular since the latter are often used for routing applications for which similar connectivity is desirable. The algorithm by Karagiorgou and Pfoser (27) computes shortest paths in both graphs with nearby start and end points, and then uses the Discrete Fréchet distance and the Average Vertical distance to compare the paths. Ahmed and Wenk (4) cover the graphs with paths of link-length three and map-match the paths to the other graph using the Fréchet distance.

3.2 Quality Measures used for Comparison

Here we give some more details on the quality measures that we use in Section 5 to compare the different map construction algorithms. Note that in our experiments the ground-truth G is an unmodified street map from OpenStreetMap and thus expected to be a superset of the underlying graph. We use the Directed Hausdorff distance (6), the path-based distance measure presented by Ahmed et al. (3), the distance measure based on shortest paths by Karagiorgou and Pfoser (27) and graph-sampling based distance measure by Biagioni and Eriksson (7). The first two measures have not been used for comparative evaluations of map constructions before.

3.2.1 Directed Hausdorff Distance (6)

The *Directed Hausdorff distance* of two sets of points A, B is defined as $\vec{d}(A, B) = \max_{a \in A} \min_{b \in B} d(a, b)$. Here, $d(a, b)$ is usually the Euclidean distance between two points a and b . In order to compare two graphs, we identify each graph as the set of points that is covered by all its vertices and edges. If the Directed Hausdorff distance from graph C to graph G is at most ε , this means that for every point on (any edge or vertex of) C there is a point on G at distance at most ε . This distance measure gives a notion about spatial distance for graphs. If C is the constructed graph and G is the ground-truth, the lower the distance from C to G , the closer the graph C to G . For computation purposes, one can consider C as set of line segments and determine how much G needs to be fattened such that each line segment of C is contained in the fattened region of G .

3.2.2 Path-Based Distance (3)

The path-based map distance considers graphs as sets of paths. The distance between two sets of paths is then computed in the Hausdorff setting, while the Fréchet distance which is a natural distance measure for curves that takes monotonicity and continuity into account, is used to compute the distance between two paths.

For curves f, g , the Fréchet distance is defined as

$$\delta_F(f, g) = \inf_{\alpha, \beta: [0,1] \rightarrow [0,1]} \max_{t \in [0,1]} d(f(\alpha(t)), g(\beta(t))), \quad (1)$$

where α, β range over continuous, surjective and non decreasing reparametrizations.

A common intuition is to explain it as the minimum leash length required such that a man and dog can walk on the two curves from beginning to end in a monotonic way.

Under this scope, let C and G be two planar geometric graphs, and let π_C be a set of paths generated from C , and π_G be a set of paths generated from G . The *path-based distance* is defined as:

$$\vec{d}_{C,G}(\pi_C, \pi_G) = \max_{p_C \in \pi_C} \min_{p_G \in \pi_G} \delta_F(p_C, p_G) \quad (2)$$

Ideally, π_C and π_G should be the set of all paths in C and G , which however has exponential size. In (3) they showed that $\vec{d}_{C,G}(\Pi_C, \Pi_G)$ can be approximated using $\vec{d}_{C,G}(\Pi_C^3, \Pi_G)$ in polynomial time using the map-matching algorithm of

(5), under some assumptions on C . Here, Π_C is the set of all paths and Π_C^3 is the set of all link-3 paths of C . A link- k path consists of k “edges”, where vertices of degree two in the graph are not counted as vertices. Using this asymmetric distance measure $\vec{d}_{C,G}(\Pi_C^k, \Pi_G)$, which can be computed in polynomial time for constant k , the following properties have been shown in (3), under some assumptions on C :

1. $k = 1$: For each edge in C , there is a path in G which is within Fréchet distance $\vec{d}_{C,G}(\Pi_C^1, \Pi_G)$. 2. $k = 2$: For each vertex v in C there is a vertex in G within bounded distance $\vec{d}_{C,G}(\Pi_C^2, \Pi_G) / \sin \frac{\theta}{2}$, where θ is the minimum incident angle at v between its adjacent edges. 3. $k = 3$: $\vec{d}_{C,G}(\Pi_C^3, \Pi_G)$ approximates $\vec{d}_{C,G}(\Pi_C, \Pi_G)$ within a factor of $1 / \sin \frac{\theta}{2}$ if the vertices of C are reasonably well separated and have degree $\neq 3$.¹

Similar to Directed Hausdorff distance, the lower the value of $\vec{d}_{C,G}(\pi_C, \pi_G)$ the more closely the constructed map C resembles the ground-truth map G .

3.2.3 Shortest-Path Based Distance (27)

Karagiorgou et al. (27) propose a measure that essentially samples each graph using random sets of shortest paths. Given the constructed and ground-truth networks C and G respectively, a common set of node pairs (origin, destination) is selected in both, using the nearest neighbor search if necessary. For all node pairs, shortest paths are computed in both networks. The geometric difference/similarity between the respective shortest paths is used to assess the similarity between C and G and consequently as a means to assess the quality of the constructed network. The Discrete Fréchet distance and the Average Vertical distance are used to compare the shortest paths. The rationale for using this approach is that measuring the similarity for sets of paths instead of individual links allows one to better reason about the connectivity of the generated network. The more “similar” the shortest paths in the constructed network are to the ground-truth network, the higher also the quality of the network. The results of this shortest-path based distance measure can be assessed by plotting the distance of all paths against each other, or by comparing average values for the entire set of paths. We employ both approaches in our experiments below.

3.2.4 Graph-Sampling Based Distance (7)

Biagioni and Eriksson (7) introduce a graph-sampling based distance measure in order to evaluate geometry and topology of the maps. The main idea is as follows: starting from a random street location, explore and place point samples on each map outward within a maximum radius. This produces two sets of locations, which are essentially spatial samples of a local map neighborhood. These two point sets are compared using one-to-one bottleneck matching and counting the unmatched points in each set. The sampling process is repeated for several seed locations.

For the bottleneck matching, the sample points on one graph can be considered as “marbles” and on the other graph as “holes”. Intuitively, if a marble lands close to a hole, it falls in, marbles that are too far from a hole re-

¹ The degree assumption is only a technical requirement for the theoretical quality guarantees, and the authors have shown (3) that similar approximation guarantees appear to hold in practice as well.

main where they land, and holes with no marbles nearby remain empty. If one of the maps is the ground truth, this difference represents the accuracy of the other map. Counting the number of unmatched marbles and empty holes quantifies the accuracy of the generated map with respect to the ground truth according to two metrics. The first metric is the proportion of spurious marbles, $spurious = spurious_marbles / (spurious_marbles + matched_marbles)$ and the second is the proportion of missing locations (empty holes), where $missing = empty_holes / (empty_holes + matched_holes)$.

To produce a combined performance measure from these two values, the well-known F-score is used, which is computed as follows:

$$F\text{-score} = 2 * \frac{precision * recall}{precision + recall} \quad (3)$$

where, $precision = 1 - spurious$ and $recall = 1 - missing$.

The higher the F-score, the closer the match. Sampling the maps locally is an important aspect of this approach as it provides the ability to capture the connectivity of the maps at a very detailed level, allowing the topological similarity to be measured. Repeated local sampling at randomly chosen locations yields an accurate view of local geometry and topology throughout the map.

A modified version is used in (8) where the method ignores parts of the map where no correspondence could be found between generated and ground-truth maps, for our experiments we used this modified version.

4 Datasets

A basic means for assessing map construction algorithms is the underlying dataset comprising vehicle trajectories and ground-truth road network datasets. The datasets are in a projected coordinate system (UTM, GGRS87). All the visualizations of the datasets are also available on the mapconstruction.org web site. The statistics of the datasets are provided in Table 2.

4.1 Tracking Data

Our experiments use several tracking datasets from different cities (Figure 2) produced by different types of vehicles, at varying sampling rates and representing different network sizes. The *Athens large* dataset contains 120 tracks with a total length of 6,781km (average: 13.27km and standard deviation: 10.79km) obtained from school buses covering an area of 12km × 14km; the tracks range from 32 to 80 position samples, with a sampling rate of 20s to 30s (average: 30.14s and standard deviation: 24.77s) and an average speed of 20.16km/h. The *Athens small* dataset contains 129 tracks with a total length of 443km (average: 3.82km and standard deviation: 1.45km) obtained from school buses covering an area of 2.6km × 6km; the tracks range from 13 to 47 position samples, with a sampling rate of 20s to 30s (average: 34.07s and standard deviation: 31.92s) and an average speed of 19.55km/h. The *Berlin* dataset contains 26,831 tracks with a total length of 41,116km (average: 1.53km and standard deviation: 634.51m) obtained from a taxi fleet covering an area of 6km × 6km; the tracks range from 22 to 58 position

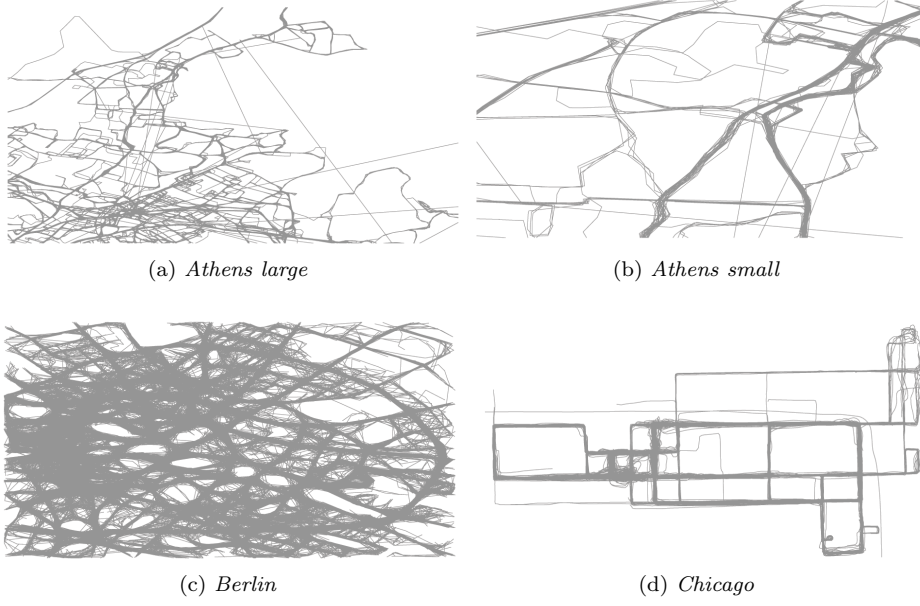


Fig. 2: Tracking data.

samples, with a sampling rate of $15s$ to $127s$ (average: $41.98s$ and standard deviation: $38.70s$) and an average speed of $35.23km/h$. It is $20\times$ the size of the *Chicago* dataset.

The *Chicago* dataset (7; 8) contains 889 tracks with a total length of $2,869km$ (average: $3.22km$ and standard deviation: $894.28m$) obtained from university shuttle buses covering an area of $7km \times 4.5km$; the tracks range from 100 to 363 position samples, with a sampling rate of $1s$ to $29s$ (average: $3.61s$ and standard deviation: $3.67s$) and an average speed of $33.14km/h$.

4.2 Ground-Truth Map Data

In all cases, we consider as ground-truth map data the corresponding road network obtained from OSM.

In *Athens large*, the road network consists of 39,699 edges and 32,212 vertices. It covers an area of $12km \times 14km$. The edges have a length of $2,000km$. The *Athens small* dataset, the road network consists of 3,436 edges and 2694 vertices. It covers an area of $2.6km \times 6km$. The edges have a length of $193km$. In *Berlin*, the road network consists of 6,839 edges and 5,894 vertices. It covers an area of $6km \times 6km$. The edges have a length of $360km$. For *Chicago*, the road network covered by the trajectories consists of 11,801 edges and 9,429 vertices. It covers an area of $7km \times 4.5km$. The edges have a length of $61km$.

Tracking Data	Trajectories	Sampling rate (s)	Trajectory length (km)	Speed (km/h)
<i>Athens large</i>	120	30.14	6,781	20.16
<i>Athens small</i>	129	34.07	443	19.55
<i>Berlin</i>	26,831	41.98	41,116	35.23
<i>Chicago</i>	889	3.61	2869	33.14

OSM Network	Vertices	Edges	Length (km)	Area (km ²)
<i>Athens large</i>	32,212	39,699	2,000	12 × 14
<i>Athens small</i>	2,694	3,436	193	2.6 × 6
<i>Berlin</i>	5,894	6,839	360	6 × 6
<i>Chicago</i>	9,429	11,801	61	7 × 4.5

Table 2: Statistics for datasets used.

5 Experiments

What follows is a description of the map construction experiments that were conducted for the range of algorithms, datasets and evaluation measures. Not all algorithms were tested on all datasets since some of them experienced either out-of-memory problems or other execution problems when tested on large datasets. Results for all algorithms are available for the small datasets, i.e., *Athens small* and *Chicago*.

Figure 3 illustrates the ground-truth road network (light grey) and the generated road networks (black) for the small *Chicago* dataset. On larger datasets, i.e., *Athens large* and *Berlin*, we ran the algorithms described in Subsections 2.2.1, 2.2.6 and 2.2.7. Figure 4 illustrates the ground-truth road network (light grey) and the generated road networks (black) for the case of larger *Berlin* dataset.

Each of the algorithms uses different parameter settings. For Ahmed and Wenk (4) the values of ε to cluster subtrajectories are: 180, 90, 170 and 80 meters for *Athens large*, *Athens small*, *Berlin* and *Chicago*, respectively. The respective parameters of *proximity* and *bearing* for the other algorithms are Biagioni 50m (8), Cao 20m and 45° (11), Davies 16m (16) and Edelkamp 50m and 45° (17). For Karagiorgou and Pfoser (27) the values of *direction*, *speed* and *proximity* to extract intersection nodes and to merge trajectories into links are 15°, 40km/h and 25m accordingly. We evaluated all constructed maps using the distance measures described in Subsection 3.2.

A summary of the complexities of the constructed maps is shown in Table 3. Here, the number of vertices includes vertices of degree two (which may lie on a polygonal curve describing a single edge), the number of edges refers to the number of undirected line segments between these vertices, and the total length refers to the total length of all undirected line segments. It appears that the *point clustering algorithms based on kernel density estimation* such as Biagioni et al. (7; 8) and Davies et al. (16) produce maps with lower complexity (fewer number of vertices and edges) but often *fail to reconstruct streets that are not traversed frequently enough* by the input tracks. In particular, the maps reconstructed by Davies et al.’s algorithm are very small. On the other hand, the algorithm by Ge et al. (22) subsample all tracks to create a much denser output set, hence the complexity of their constructed maps is always higher.

Map construction algorithms based on *incremental track insertion*, such as Ahmed et al. (4) and Cao et al. (11) fail to cluster tracks together when the variability

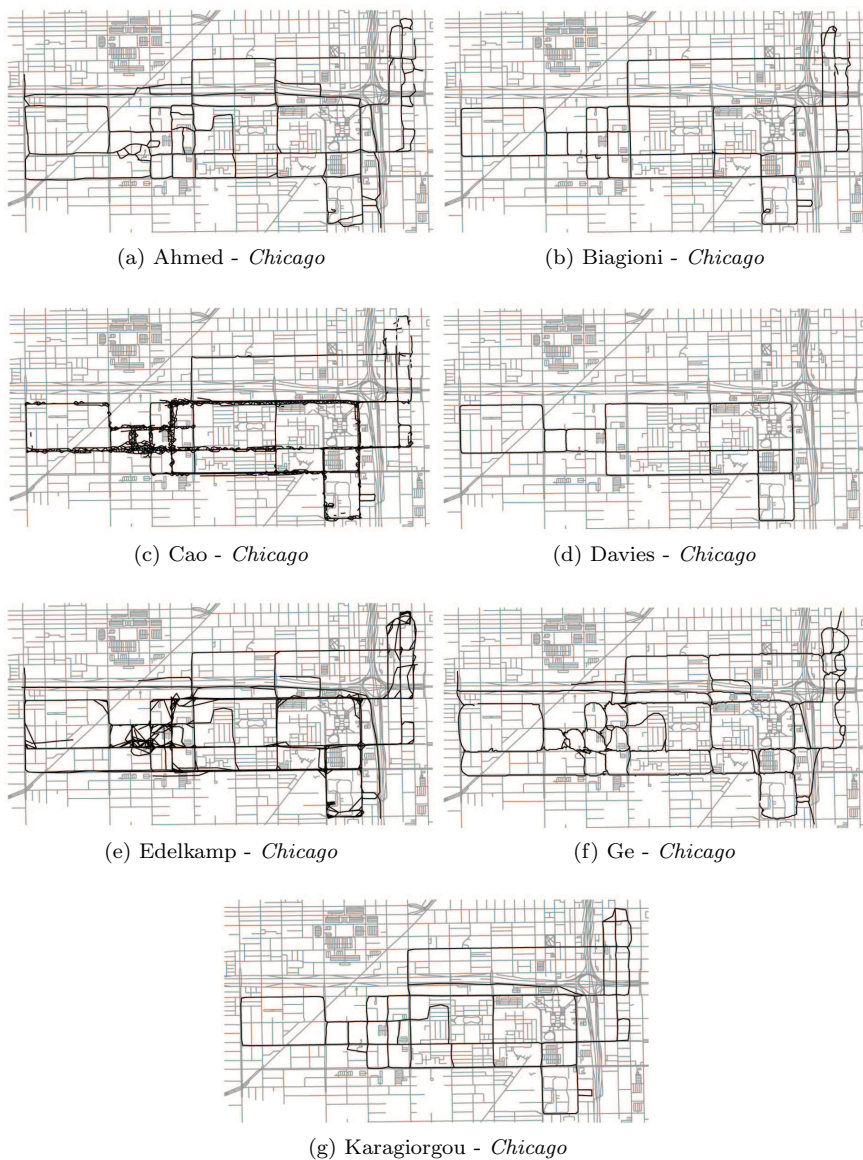


Fig. 3: Generated road networks (small dataset).

and error associated with the input tracks is large. As a result, the constructed street maps contain *multiple edges for a single street*, which implies larger values in the total edge length column in Table 3.

Several examples of generated road networks are shown in Figure 3 and Figure 4. Since not all algorithms produced results for all networks, we showcase examples of the smaller *Chicago* network in Figure 3. It can be clearly seen that the




Fig. 4: Generated road networks (large dataset).

Generated Road Network	# Vertices	# Edges	Length (km)
<i>Athens large</i>			
Ahmed	7067	7960	1358
Ge	20774	21626	9740
Karagiorgou	6584	5280	252
<i>Athens small</i>			
Ahmed	344	378	35
Biagioni	391	398	22
Cao	20	14	3
Davies	209	227	2
Edelkamp	526	1037	197
Ge	1936	1993	23
Karagiorgou	660	637	35
<i>Berlin</i>			
Ahmed	1322	1567	164
Ge	15450	16136	183
Karagiorgou	2542	2262	161
<i>Chicago</i>			
Ahmed	1195	1286	34
Biagioni	303	322	24
Cao	2092	2948	78
Davies	1277	1310	14
Edelkamp	828	1247	83
Ge	5893	6672	37
Karagiorgou	596	558	26

Table 3: Generated road networks complexities.

coverage and quality of the constructed road network varies considerably. Three examples for the *Berlin* network are also given in Figure 4. More examples can be found on the mapconstruction.org web site.

5.1 Path-Based and Hausdorff Distance

For the *path-based distance measure* we generated all paths of link-length 3 for each generated map. For each path, we computed the Fréchet distance between the path and the ground-truth map. We then computed the minimum, maximum, median, average of all the obtained distances. We also computed the $d\%$ -distance, as the maximum of the distances after removing the $d\%$ largest distances (“outliers”). For the *Directed Hausdorff distance*, we computed all link-length 1 paths and computed the Directed Hausdorff distance of the union of all edges to the ground-truth map. Our results are summarized in Table 4. In the case of *Athens small*, the Cao algorithm produced a very small road network and thus it was not possible to perform a quantitative evaluation.

The maps reconstructed using the algorithms by Karagiorgou et al. (27) and by Biagioni et al. (7; 8) generally have a better path-based distance than the others. Note that Davies et al.’s (16) map is unusually small for the *Athens small* dataset. Their ideas of averaging trajectories or computing skeletons seem to help to improve the quality of the edges of the produced map.

Generated Road Network	Path based distance (m)								Directed Hausdorff distance (m)							
	min	max	median	avg	2%	5%	10%	15%	min	max	median	avg	2%	5%	10%	15%
<i>Athens large</i>																
Ahmed	7	849	70	85	250	164	132	114	1	269	30	33	84	67	56	50
Ge	7	956	76	90	237	188	150	116	1	295	35	37	95	74	59	52
Karagiorgou	2	175	25	32	109	80	63	53	1	200	10	13	46	35	26	22
<i>Athens small</i>	min	max	median	avg	2%	5%	10%	15%	min	max	median	avg	2%	5%	10%	15%
Ahmed	9	224	45	52	101	101	81	72	1	82	25	26	82	54	46	40
Biagioni	5	73	35	36	67	66	61	57	3	74	19	20	47	43	31	31
Cao									5	25	13	13	25	25	25	22
Davies	4	38	11	11	38	18	14	14	2	13	7	6	13	13	13	11
Edelkamp	2	229	36	39	89	72	68	61	1	86	18	21	63	50	42	37
Ge	19	251	52	59	142	113	89	76	3	81	21	23	80	59	39	35
Karagiorgou	7	229	32	38	113	68	59	57	2	84	14	17	54	40	33	30
<i>Berlin</i>	min	max	median	avg	2%	5%	10%	15%	min	max	median	avg	2%	5%	10%	15%
Ahmed	9	540	66	74	207	147	120	107	1	219	30	33	95	70	60	53
Ge	13	808	65	75	214	157	117	103	4	562	36	37	73	62	55	51
Karagiorgou	4	306	28	37	120	85	65	52	1	232	14	18	59	42	34	30
<i>Chicago</i>	min	max	median	avg	2%	5%	10%	15%	min	max	median	avg	2%	5%	10%	15%
Ahmed	7	201	35	42	127	100	85	76	1	81	14	19	72	59	43	35
Biagioni	3	71	15	18	71	38	27	26	2	53	9	11	29	25	23	17
Cao	1	126	24	27	79	61	49	42	1	78	9	12	44	35	28	25
Davies	2	92	12	14	57	24	22	21	2	20	8	7	20	14	13	12
Edelkamp	1	205	29	37	99	84	72	66	1	93	8	13	57	48	35	25
Ge	18	346	50	56	158	126	95	75	7	72	26	28	64	61	53	46
Karagiorgou	3	89	15	23	72	72	65	51	1	48	7	8	41	23	15	13

Table 4: Path-Based and Directed Hausdorff distance measure evaluation.

For further analysis of the results, we selected the *Chicago* dataset as all map construction algorithms produced results for it. From Table 4 one can see that the path-based distance and the Directed Hausdorff distance are smaller for the generated maps by Biagioni, Davies and Karagiorgou (shaded grey) compared to map generated using other algorithms. A visual inspection of the maps in Figure 3 justifies the result. Note that Davies et al.’s (16) map is comparatively smaller than the other, see Table 3. Although, the algorithms by Ahmed et al. and by Ge et al. produce maps with good coverage, their path-based distances are larger

since they employ less aggressive averaging techniques that would help cope with noise in the input tracks.

To illustrate the appropriateness of the path-based distance, consider the path in Figure 5 from the map generated by Biagioni et al. This is an example where the Fréchet-based distance measure is more effective than any point-based measure. As Fréchet distance ensures continuous mapping, the whole path needs to be matched with the bottom horizontal edge of the ground-truth map. The Fréchet distance for this path is $71m$. For the same path, the Hausdorff distance is $53m$, as this only requires for each point on the path to have a point on the graph close-by. So, to evaluate the connectivity of a map, the Fréchet distance is a more suitable distance measure than any point-based measure.




Fig. 5: A path with Fréchet distance greater than Hausdorff distance.






Fig. 6: Distributions of individual path distances (Biagioni alg. - *Chicago*).

In addition, if desired one can discard outliers by computing the $d\%$ -distance. Figure 6 shows the distribution of both the path-based measure and the Directed Hausdorff distance for Biagioni et al. In both cases, a very small number of paths have the maximum distance, and the distances for most of the paths are distributed within a small range. Removing only 5% of the outliers brings the path-based distance from $71m$ to $38m$ and the Directed Hausdorff distance from $53m$ to $25m$. Figure 7 shows paths with larger distances and with smaller distances in different colors. Such visual representation helps to identify areas in the map that have higher distance to the ground-truth map.



(a) Biagioni et al. Blue paths have Fréchet distance < 15 (median of all values) and red paths ≥ 15



(b) Ahmed et al. Blue paths have Fréchet distance < 51 (median of all values) and red paths ≥ 51

Fig. 7: All link-length 3 paths, color indicates higher or lower Fréchet distance than median.

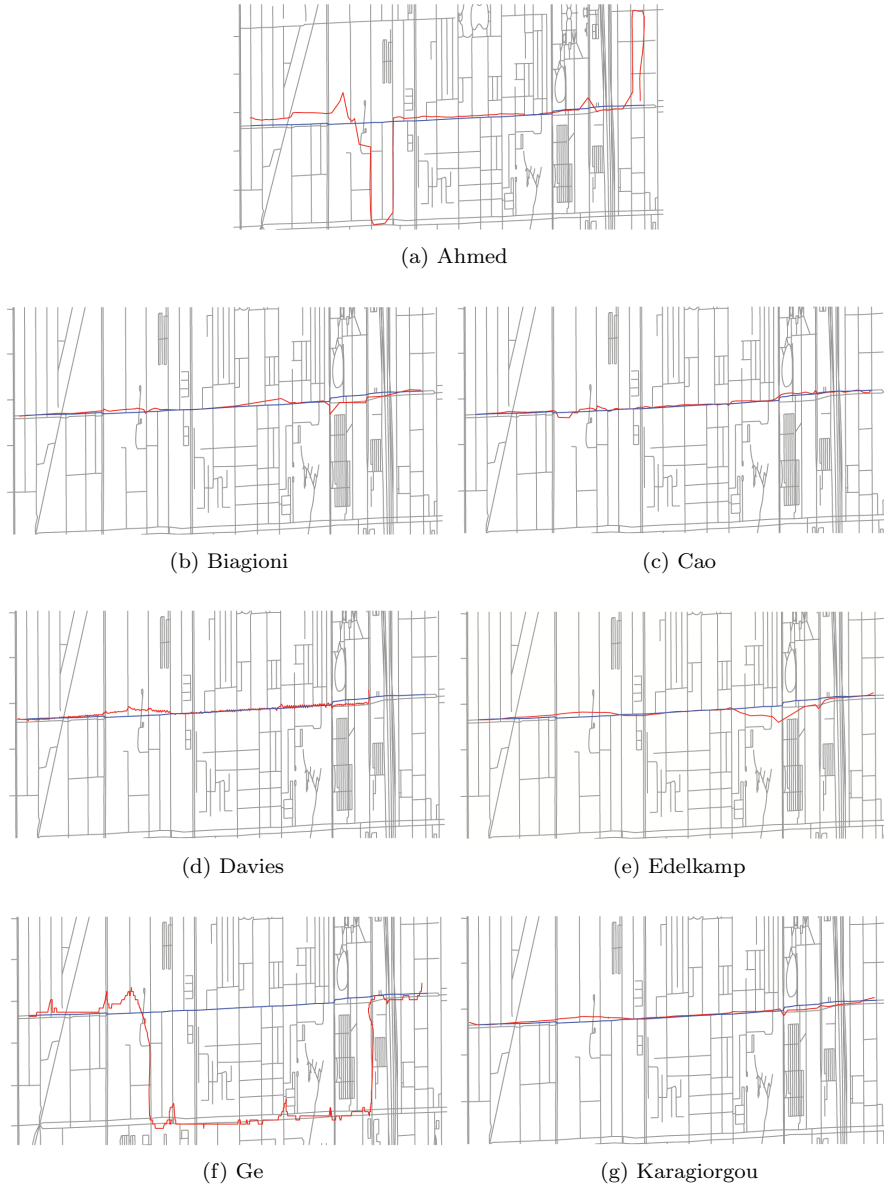


Fig. 8: Examples of shortest paths for the *Chicago* dataset.

5.2 Shortest-Path Based Measure

Another means to compare the constructed networks is the shortest-path based distance. For each city, we computed a set of 500 random shortest-paths with origin

and destination nodes uniformly distributed over the networks and compared the paths using the Discrete Fréchet and Average Vertical distance measure.

A first impression on how different constructed road networks affect such paths is given in Figure 8. Given a specific origin and destination for the *Chicago* road network, the shortest path has length $3.66km$ in the ground-truth network. In the map generated by Ahmed et al.’s algorithm the shortest path has length $4.67km$ (a Discrete Fréchet distance with respect to the ground-truth network of $65m$, and an Average Vertical distance of $21m$). The respective results for the other algorithms are Biagioni $3.71km$ ($36m$, $5m$), Cao $3.76km$ ($24m$, $6m$), Davies $3.39km$ ($35m$, $4m$), Edelkamp $3.64km$, ($26m$, $8m$), Ge $7.33km$, ($174m$, $98m$), and Karagiorgou $3.73km$ ($21m$, $5m$). For most algorithms the resulting paths have small distance to the shortest path in the ground-truth network. However, in the case of Ahmed (Figure 8a) and Ge (Figure 8f), due to significant differences in the generated networks, different shortest paths have been computed that have a larger distance to the shortest path in the ground-truth network. This result is in line with the path-based measure of Section 5.1, where also Biagioni, Davies and Karagiorgou produced the best constructed maps.

Figures 9a and 9b show the Discrete Fréchet and the Average Vertical distance measures for each of the 500 paths per algorithm for the *Athens large* network. The paths are ordered by increasing distance of the shortest-path length with respect to the ground-truth road network. Some paths could not be computed for some networks due to connectivity problems (missing links). Some other paths experience greater distance measures due to spatial accuracy problems. The graph shows that some algorithms produce networks, which resemble the actual road network more closely as assessed by this shortest-path sampling approach.

Finally, the shortest-path based evaluation is summarized in Table 5. The first column shows the percentage (%) of shortest-paths that in each case could be computed, i.e., an algorithm might find an accurate, but small network. The second and the third column show the two different distance measures used to compare the resulting paths. The fourth column gives some statistics with respect to the computed shortest paths. Considering the example of *Berlin* and here the Ahmed algorithm result (shaded light gray) in Table 5, this algorithm produces a network that in turn generates paths that have a min, max, and avg. Discrete Fréchet distance of $21m$, $469m$, and $192m$, respectively. An aspect not captured by these distances are missing paths due to *limited network coverage*. Consider the case of Davies for *Chicago* and Cao for *Athens small* (shaded light gray in Table 5). In both cases, the distance measures suggest good network quality. However, in both cases the constructed network has a small coverage, as only 92.6% and 7.0% of the 500 total paths were computed. In this evaluation, Karagiorgou produces networks that have, both, good coverage and high path similarity (cf. dark-shaded entry for Berlin - good coverage and small distance measure indicating similar paths between constructed and ground-truth network).

Overall, shortest-path sampling provides an effective means for assessing the quality of constructed road networks as it not only considers *similarity*, but also the *coverage of the road network*.



(a) Discrete Fréchet distance - *Athens large*(b) Average Vertical distance - *Athens large*

Fig. 9: Road networks comparison.

Generated Road Network	Found (%)	Discrete Fréchet dist. (m)				Average Vertical dist. (m)				Shortest-path dist. (km)			
<i>Athens large</i>		min	max	avg	stddev	min	max	avg	stddev	min	max	avg	stddev
Ahmed	92.6	23	445	137	103	12	230	106	62	1.12	11.84	6.93	2.92
Ge	92.8	25	497	149	112	14	241	120	65	1.47	11.91	7.13	3.18
Karagiorgou	94.2	19	432	125	96	9	225	98	58	1.01	11.62	6.84	2.86
<i>Athens small</i>		min	max	avg	stddev	min	max	avg	stddev	min	max	avg	stddev
Ahmed	97.6	13	234	96	62	6	91	38	24	1.28	5.72	3.11	1.84
Biagioni	94.2	7	214	84	50	4	80	28	21	0.79	5.23	2.97	1.41
Cao	7.0	7	26	10	11	4	13	6	5	0.17	0.31	0.22	0.21
Davies	22.6	9	258	102	69	5	81	31	22	0.85	5.25	2.99	1.47
Edelkamp	97.2	15	228	97	64	6	93	40	26	0.93	5.29	3.02	1.51
Ge	93.4	21	290	123	75	11	127	63	33	1.43	5.93	3.41	1.92
Karagiorgou	96.8	7	212	81	48	3	81	27	20	0.78	5.21	2.95	1.39
<i>Berlin</i>		min	max	avg	stddev	min	max	avg	stddev	min	max	avg	stddev
Ahmed	93.2	21	469	191	123	12	231	121	63	1.56	5.88	3.49	1.96
Ge	92.4	25	475	194	128	15	236	127	64	1.85	5.93	3.84	2.03
Karagiorgou	93.8	18	428	183	112	8	209	106	58	1.32	5.67	3.27	1.84
<i>Chicago</i>		min	max	avg	stddev	min	max	avg	stddev	min	max	avg	stddev
Ahmed	99.8	13	208	97	56	6	92	43	19	1.21	6.95	4.45	2.04
Biagioni	98.6	4	98	40	27	2	49	20	13	0.89	6.03	3.76	1.57
Cao	99.2	7	131	67	34	4	76	41	17	1.02	6.87	3.94	1.84
Davies	92.6	5	97	41	27	3	51	23	15	0.93	6.08	3.88	1.66
Edelkamp	99.0	12	211	98	58	5	89	41	18	1.19	6.88	4.32	1.97
Ge	99.8	19	241	127	63	8	94	49	22	1.58	6.98	4.69	2.25
Karagiorgou	99.2	4	103	41	28	2	50	21	14	0.90	6.05	3.82	1.59

Table 5: Shortest-path measure evaluation summary.

5.3 Graph-Sampling Based Distance

For this measure we use the source code obtained from the authors of (7). We modified the code to use Euclidean distance as our data uses projected coordinate system. The method that computes this measure has four parameters: 1. *sampling density*, how densely the map should be sampled (marbles for generated map and holes for ground-truth map), we use 5 meters. 2. *matched distance*, the maximum distance between a matched marble-hole pair, we vary this distance from 10 to 120m. 3. *maximum distance from root*, the maximum distance from randomly selected start location one will explore, we use 300m. 4. *number of runs*, number of start locations to consider, we use 1000. To make our comparison of all generated maps consistent, we generated a sequence of random locations for each dataset and used the first 1,000 locations from the same sequence for each algorithm for which both maps (ground-truth and generated) had correspondences within *matched distance*. When two maps are very similar, they should have very few unmatched marbles and holes, which implies the precision, recall and F-score values should be very close to 1. In our case, as we used a superset of the ground-truth map, there should be a large number of unmatched holes, which implies smaller recall and F-score values, but still the relative comparison of F-score values should provide an idea if an algorithm performs better than another one.




Fig. 10: Comparison of F-scores - *Chicago*.

Figure 10 shows F-score values for the *Chicago* dataset for different generated maps. As our ground-truth is a superset of the actual ground-truth represented by the tracking dataset, a larger matching distance creates unexpected results for algorithms that generate extra edges and vertices. For example, Cao and Edelkamp for *Chicago*, the precision is low as there will be lots of unmatched marbles (cf. entry for Cao and Edelkamp for *Chicago* in Table 6). However, a larger matching distance decreases the number of unmatched marbles by matching these with available holes

Generated Road Network	Precision Value (for <i>matched distance</i> 10, 40, 70, 100)			
<i>Athens large</i>	10	40	70	100
Ahmed	0.216	0.407	0.497	0.591
Ge	0.149	0.368	0.507	0.635
Karagiorgou	0.394	0.559	0.630	0.711
<i>Athens small</i>	10	40	70	100
Ahmed	0.265	0.442	0.503	0.579
Biagioni	0.450	0.586	0.662	0.727
Cao	0.415	.691	0.722	0.810
Davies	0.439	0.574	0.617	0.670
Edelkamp	0.106	0.156	0.197	0.232
Ge	0.409	0.527	0.624	0.708
Karagiorgou	0.343	0.489	0.561	0.647
<i>Berlin</i>	10	40	70	100
Ahmed	0.123	0.326	0.422	0.485
Ge	0.142	0.457	0.534	0.584
Karagiorgou	0.294	0.590	0.633	0.649
<i>Chicago</i>	10	40	70	100
Ahmed	0.312	0.563	0.658	0.738
Biagioni	0.491	0.699	0.730	0.775
Cao	0.209	0.321	0.376	0.456
Davies	0.488	0.650	0.690	0.739
Edelkamp	0.334	0.431	0.473	0.541
Ge	0.306	0.487	0.565	0.645
Karagiorgou	0.602	0.740	0.751	0.801

Table 6: Precisions for varying *matched distance*.

that probably are not part of the actual ground-truth. A higher recall value yields a higher F-score, which does not necessarily reflect better-quality maps (cf. Figure 3 and Figure 4).

Hence, in Table 6 we are ignoring F-score and recall values and *showcase only precision values*. According to precision values, the algorithms by Biagioni, Davies and Karagiorgou perform best for dataset *Chicago*, which is consistent with our findings using the other three distance measures.

6 Conclusions

This survey has considered the active field of map construction and has considered a variety of map construction algorithms. In the past, the lack of benchmark data and quantitative evaluation methods has hindered a cross-comparison between algorithm. In this paper, the contribution of benchmark data sets and code for map construction algorithms and evaluation measures for the first time enables a standardized assessment and comparison of map construction algorithms. All data, map construction, and evaluation algorithms are available with detailed execution instructions on the mapconstruction.org web site. Directions for future work are the expansion of the web site towards the inclusion of more algorithms and source code. The final goal will be to provide an easy-to-use benchmark suite and automated quality measurements for generated maps.

Acknowledgements This work has been supported by the National Science Foundation grant CCF-1301911 and the European Union Seventh Framework Programme - Marie Curie Actions, Initial Training Network GEOCROWD (www.geocrowd.eu) under grant agreement No. FP7-PEOPLE-2010-ITN-264994. We thank Yusu Wang and Xiaoyin Ge for running their

map construction algorithm on our benchmark datasets and James Bagioni for making the source code of several algorithms for map construction and map compare available to us. Associated data and software will be made available at mapconstruction.org.

References

1. Aanjaneya, M., Chazal, F., Chen, D., Glisse, M., Guibas, L.J., Morozov, D.: Metric graph reconstruction from noisy data. In: Proc. 27th ACM Symp. on Comp. Geometry, pp. 37–46 (2011)
2. Agamennoni, G., Nieto, J.I., Nebot, E.M.: Robust inference of principal road paths for intelligent transportation systems. *Trans. Intell. Transport. Sys.* **12**(1), 298–308 (2011)
3. Ahmed, M., Hickmann, K.S., Wenk, C.: Path-based distance for street map comparison. *arXiv:1309.6131* (2013)
4. Ahmed, M., Wenk, C.: Constructing street networks from GPS trajectories. In: Proc. 20th Annual European Symposium on Algorithms, pp. 60–71 (2012)
5. Alt, H., Efrat, A., Rote, G., Wenk, C.: Matching planar maps. *Journal of Algorithms* pp. 262–283 (2003)
6. Alt, H., Guibas, L.: Discrete geometric shapes: Matching, interpolation, and approximation - a survey. In: J.R. Sack, J. Urrutia (eds.) *Handbook of Computational Geometry*, pp. 121–154. Elsevier (1999)
7. Biagioni, J., Eriksson, J.: Inferring road maps from global positioning system traces: Survey and comparative evaluation. *Transportation Research Record: Journal of the Transportation Research Board* **2291**, 61–71 (2012)
8. Biagioni, J., Eriksson, J.: Map inference in the face of noise and disparity. In: Proc. 20th ACM SIGSPATIAL, pp. 79–88 (2012)
9. Brakatsoulas, S., Pfoser, D., Salas, R., Wenk, C.: On map-matching vehicle tracking data. In: Proc. 31st VLDB Conference, pp. 853–864 (2005)
10. Bruntrup, R., Edelkamp, S., Jabbar, S., Scholz, B.: Incremental map generation with GPS traces. In: Proc. IEEE Intelligent Transp. Sys., pp. 574 – 579 (2005)
11. Cao, L., Krumm, J.: From GPS traces to a routable road map. In: Proc. 17th ACM SIGSPATIAL, pp. 3–12 (2009)
12. Chen, C., Cheng, Y.: Roads digital map generation with multi-track GPS data. In: Proc. Workshops on Education Technology and Training, and on Geoscience and Remote Sensing, pp. 508–511. IEEE (2008)
13. Chen, D., Guibas, L.J., Hershberger, J., Sun, J.: Road network reconstruction for organizing paths. In: Proc. 21st ACM-SIAM Symp. on Discrete Algorithms, pp. 1309–1320 (2010)
14. Cheong, O., Gudmundsson, J., Kim, H.S., Schymura, D., Stehn, F.: Measuring the similarity of geometric graphs. In: *Proceedings of the 8th International Symposium on Experimental Algorithms*, SEA '09, pp. 101–112. Springer-Verlag, Berlin, Heidelberg (2009)
15. Conte, D., Foggia, P., Sansone, C., Vento, M.: Thirty years of graph matching in pattern recognition. *International Journal of Pattern Recognition and Artificial Intelligence* **18**(3), 265–298 (2004)
16. Davies, J.J., Beresford, A.R., Hopper, A.: Scalable, distributed, real-time map generation. *IEEE Pervasive Computing* **5**(4), 47–54 (2006)
17. Edelkamp, S., Schrödl, S.: Route planning and map inference with global positioning traces. In: *Computer Science in Perspective*, pp. 128–151. Springer (2003)
18. Efentakis, A., Brakatsoulas, S., Grivas, N., Lamprianidis, G., Patroumpas, K., Pfoser, D.: Towards a flexible and scalable fleet management service. In: Proc. 6th ACM SIGSPATIAL Int'l Workshop on Computational Transportation Science (IWCTS), pp. 79–84 (2013)
19. Fathi, A., Krumm, J.: Detecting road intersections from GPS traces. In: Proc. 6th Int. Conf. on Geographic Information Science, pp. 56–69 (2010)
20. Gao, X., Xiao, B., Tao, D., Li, X.: A survey of graph edit distance. *Pattern Analysis and Applications* **13**, 113–129 (2010)
21. Gati, G.: Further annotated bibliography on the isomorphism disease. *Journal of Graph Theory* **3**(2), 95–109 (1979)
22. Ge, X., Safa, I., Belkin, M., Wang, Y.: Data skeletonization via Reeb graphs. In: Proc. 25th Annual Conference on Neural Information Processing Systems, pp. 837–845 (2011)

23. Goodchild, M.F.: Citizens as voluntary sensors: spatial data infrastructure in the world of web 2.0. *International Journal of Spatial Data Infrastructures Research* **2**, 24–32 (2007)
24. Guo, T., Iwamura, K., Koga, M.: Towards high accuracy road maps generation from massive GPS traces data. In: *Proc. IEEE International Geoscience and Remote Sensing Symposium*, pp. 667–670 (2007)
25. Haklay, M., Weber, P.: Openstreetmap: User-generated street maps. *IEEE Pervasive Computing* **7**(4), 12–18 (2008)
26. Jang, S., Kim, T., Lee, E.: Map generation system with lightweight GPS trace data. In: *Proc. 12th Int. Conf. on Adv. Comm. Tech.*, pp. 1489–1493 (2010)
27. Karagiorgou, S., Pfoser, D.: On vehicle tracking data-based road network generation. In: *Proc. 20th ACM SIGSPATIAL*, pp. 89–98 (2012)
28. Kégl, B., Krzyzak, A., Linder, T., Zeger, K.: Learning and design of principal curves. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(3), 281–297 (2000)
29. Liu, X., Biagioni, J., Eriksson, J., Wang, Y., Forman, G., Zhu, Y.: Mining large-scale, sparse GPS traces for map inference: comparison of approaches. In: *Proc. 18th ACM SIGKDD*, pp. 669–677 (2012)
30. Niehofer, B., Burda, R., Wietfeld, C., Bauer, F., Lueert, O.: GPS community map generation for enhanced routing methods based on trace-collection by mobile phones. In: *Proc. 1st Int. Conf. on Advances in Satellite and Space Comm.*, pp. 156–161 (2009)
31. OpenStreetMap Foundation: Bulk gpx track data (2013). URL <http://blog.osmfoundation.org/2013/04/12/bulk-gpx-track-data/>
32. OpenStreetMap Foundation: Openstreetmap: User-generated street maps (2013). URL <http://www.openstreetmap.org>
33. Qudus, M., Ochieng, W., Noland, R.: Current map-matching algorithms for transport applications: State-of-the art and future research directions. *Transportation Research Part C: Emerging Technologies* pp. 312–328 (2007)
34. Read, R.C., Corneil, D.G.: The graph isomorphism disease. *Journal of Graph Theory* **1**(4), 339–363 (1977)
35. Rogers, S., Langley, P., Wilson, C.: Mining GPS data to augment road models. In: *Proc. 5th ACM SIGKDD*, pp. 104–113 (1999)
36. Schroedl, S., Wagstaff, K., Rogers, S., Langley, P., Wilson, C.: Mining GPS traces for map refinement. *Data Min. Knowl. Discov.* **9**, 59–87 (2004)
37. Shi, W., Shen, S., Liu, Y.: Automatic generation of road network map from massive GPS vehicle trajectories. In: *Proc. 12th International IEEE Conference on Intelligent Transportation Systems*, pp. 48–53 (2009)
38. Steiner, A., Leonhardt, A.: Map generation algorithm using low frequency vehicle position data. In: *Proc. 90th Ann. Meeting of the Transportation Research Board*, pp. 1–17 (2011)
39. Wang, Y., Liu, X., Wei, H., Forman, G., Chen, C., Zhu, Y.: Crowdatlas: Self updating maps for cloud and personal use. In: *Proc. 11th Int. Conf. Mobile Systems, Applications and Services* (2013)
40. Worrall, S., Nebot, E.: Automated process for generating digitised maps through GPS data compression. In: *Proc. Australasian Conf. on Robotics and Automation* (2007)
41. Zeng, Z., Tung, A.K.H., Wang, J., Feng, J., Zhou, L.: Comparing stars: on approximating graph edit distance. In: *Proc. 35th VLDB conf.*, pp. 25–36 (2009)
42. Zhang, L., Thiemann, F., Sester, M.: Integration of GPS traces with road map. In: *Proc. 3rd ACM SIGSPATIAL International Workshop on Computational Transportation Science*, pp. 17–22 (2010)